

Cost-Effective Integration of MKM Semantic Services into Editing Environments

Constantin Jucovschi

Jacobs University Bremen, Germany

Abstract. Integration of MKM services into editors has been of big interest in both formal as well as informal areas of MKM. Until now, most of the efforts to integrate MKM services into editing environments are done on an individual basis which results in high creation and maintenance costs.

In this paper, I propose an architecture which allows editing environments and MKM services to be integrated in a more efficient way. This is accomplished by integrating editors and services only once with a real-time document synchronization and service broker. Doing so simplifies the development of services, as well as editor integrations. Integrating new services into an arbitrary number of already integrated editors can then take as little as 3-4 hours of work.

1 Introduction

Integration of MKM services into editors has been of big interest in both formal as well as informal areas of MKM. In the formal area of MKM, it becomes more and more important to have an efficient way to work with formal documents and pass information between interactive provers and the user. Examples of useful services are: showing type information when hovering over an expression, navigating to definitions of symbols, and supporting editor features like folding inside long formulae.

Informal MKM needs editing support to make it easier for authors to create semantically annotated documents. This can mean integration of Natural Language Processing (NLP) services to e.g. spot mathematical terms as well as hiding (or folding) existing semantic annotations in order to provide a better reading experience.

Until now, most of the efforts to integrate MKM services into editing environments were done on an individual basis, namely, some service X was integrated into an editing environment Y . Obviously it would have been more efficient to integrate service X into all editing environments Y where this service makes sense. The problem is that developing and maintaining such integrations requires a lot of effort and hence MKM services are usually integrated with maybe one or two editing environments.

The high cost of semantic service integration has an especially negative impact on informal MKM because its users, the authors of mathematical texts,

write mathematics in a handful of different programs and operating systems. Even if one chooses to support the 5 most used editing environments, it would still be too expensive to maintain 10 semantic service integrations. Clearly, a different integration strategy is called for.

In this paper, I propose an architecture allowing tight integration of authoring services into the authoring process. It enables distributed services, running on different platforms and hardware, to be notified of changes made by the user to a certain document. Services can then react to these changes by modifying/coloring document's text as well as enriching parts of the text with extra (invisible to the user) semantic annotations. This architecture can accommodate both 1. reactive services like syntax highlighting giving user the illusion that the service runs natively in the editor as well as, 2. time-consuming services like Natural Language Processing tasks, without requiring the user to wait while text is processed.

In the next section I discuss in more detail what type of services would fit the presented framework, and compare it to other existing MKM frameworks. In section 3, I present the proposed integration architecture. To validate the architecture I implemented four services and integrated them in two editors. In section 4, I briefly describe these services, and give some high level technical details on how new services can be built. Section 5 concludes the paper.

2 Aims and Scope of Integration

The task of tool integration is a very complex and multi-faceted one. Many frameworks and technologies [Wic04] have been proposed to integrate tools, each optimizing some aspects of tool integration and yet, none of them is widely adopted. The current paper does not attempt to create a framework to integrate all possible editors with all possible services. It considers only pure text editors and integrates only services that participate in the editing/authoring process. The scope of the framework is purposefully kept relatively small so that it solves a well-defined part of editor-service integration problem and can eventually be used along with other integration frameworks.

In the next section, I would like to make more explicit the types of services that are included in the scope of the framework. In section 2.2 I analyze what integration strategies will be used to achieve optimal integration. This will later help me in section 2.3 to differentiate the framework more clearly from other existing frameworks.

2.1 Targeted Authoring Services

In this paper by “authoring service for text-based documents” (abbreviated as “authoring service” or just “service”), I mean a service which provides some added value to the process of authoring the text document by:

1. reacting to document changes and giving feedback e.g. by coloring/underlining parts of text,

2. performing changes to the document e.g. as result of an explicit request to autocomplete, use a template, or fold some part of text/formula.

The strategy to achieve editor-service integration is to require authoring services to be agnostic of advanced editor features like the ability to fold lines or embed images/MathML formulae. Services should only be allowed to assume a relatively simple document model that allows a limited set of operations (both described in section 3.2) and which are supported by most editing environments. Likewise, services should only assume a limited set of possible interactions with the user (see section 3.4).

A way to decide whether a certain service is in the scope of the presented integration framework is to analyze whether it can be realized conforming to the limited document model and interaction possibilities.

2.2 Levels of Integration

One of the classical ways used to describe and compare integrations was proposed by Wasserman [Was90] and later improved by Thomas and Nejme [TN92]. They propose 4 dimensions along which an integration can be analyzed, namely: presentation, data, control and process integration dimensions. Presentation integration accounts for the level at which tools share a common “look and feel”, mental models of interaction, and interaction paradigms. Data integration dimension analyzes how data is produced, shared and kept consistent among tools. Control integration dimension analyzes the level to which tools use each other’s services. The process integration dimension describes how well tools are aware of constraints, events and workflows taking place in the system.

Note that high or low level of integration does not reflect the quality of some integration. Low integration level in some dimension only suggests that those tools can be easily decoupled and interchanged. High integration level, on the other hand, suggests that tools connect in a deeper way and can enable features not possible otherwise. Experience suggests that, low level of integration require less maintenance costs in the long run and should be used whenever possible.

According to description of the targeted authoring services, I derived the integration levels that need to be supported in each integration dimension. Namely we need:

1. medium-high presentation integration level, because services need to be able to change the text/colors of the document as well as interact with the user using some high level interaction paradigms. While the type of changes/interactions a service can perform are limited, a service should be able to perform these actions unrestricted by other components.
2. high data integration level due to the fact all service are distributed and still need to be able to perform changes to the edited document. Hence synchronization and data consistency mechanisms are needed.
3. low control integration as the framework should only give support for integrating services with editors. Service-service integrations are outside the scope of integration.

4. low process integration mainly due to the fact that services are expected to be mostly stand-alone and the workflows should only involve an editor and a service.

2.3 Comparison to other MKM Integrations

A lot of integrations combining several tools have been developed in MKM. In the context of the current paper only several types of integrations are interesting:

1. frameworks that enable integration of services into authoring process in a consistent, service independent way,
2. integrations which allow multiple services to listen/react to document changes.

The Proof General (PG) framework along with the PGIP message protocol [ALW07] constitute the base for several popular integrations between editors (e.g. emacs, Eclipse) and interactive provers like Isabelle, Coq and HOL. PG framework differs from presented framework in the following ways:

1. low presentation integration — consisting of changing the color of text regions according to prover state, accompanied by locks of regions under processing. The later is not a typical interaction a user would expect. Integrated provers cannot directly influence the display or interact with the users. Hence only the broker component is directly integrated from a presentation integration point of view.
2. low data integration — editors and services communicate mainly by sending parts of the source text to the prover as result of the user changing regions of text. This data never gets changed by services and hence no synchronization or consistency checking is needed.
3. medium control integration — the broker uses different protocols to talk with displays and provers. Additionally, it has the role of orchestrating interaction between them.

The effort of Aspinnall et al. [ALW06] further extends PG architecture to integrate rendering processors (e.g. \LaTeX) and possibly other tools (e.g. code generators). These newly integrated components are loosely integrated by running them on a (hidden from the user) updated version of the original document. The broker-prover integration becomes tighter due to the documentation and script backflow mechanisms. Differences to the current framework are:

1. from the presentation integration perspective it extends the PG framework by two (or more) additional views of the document. These views are presented and updated in separate windows and provide the user with more focused views on the authored content. These new views do not seem to provide additional interactions to the user. In conclusion, the presentation integration is tighter compared to PG but still relatively low.
2. data integration becomes tighter between the broker and prover components as a result of the backflow mechanisms but is still relatively low. The additional complexity due to the backflow is mostly in the broker component which needs to know where in the central document to integrate data coming from the prover. Data passed around still never gets modified and hence no advanced synchronization or consistency checking is needed.

3. control integration is similar to the PG framework.

The integration of provers and editors is done quite differently in PIDE [Wen10]. Instead of fixing a protocol encapsulating all the features a prover can provide (like PGIP does), a document model is specified and the protocol to interact with that document model is fixed. This has the advantage that editors need to know much less about the provers features, and only need to provide them with changes the user made to the document. Conceptually this is very similar to the approach taken in the current work. The difference is that in the case of PIDE, the document model is shared only between two entities (the editor and the prover) and that these entities share the same running environment. The current architecture allows several authoring services to listen and change the shared document and allows them to be distributed.

3 Editor Service Integration Architecture

Creating and maintaining integrations between software programs is generally an expensive task. However, there are some well known best practices for these tasks which have been proven to help a lot in reducing both creation as well as maintenance costs. A good example is the integration between database systems and hundreds of languages and frameworks. Some of the key aspects that make such integrations possible are:

- P1. Client-server architecture** allows clients and servers to be developed and executed on arbitrarily different environments.
- P2. Stable communication protocol** on the server side reduces maintenance costs and makes documentation more stable and complete.
- P3. Declarative API** is usually more stable as it requires definition of a small set of primitives and some way to compose them.

The goal of the current architecture is to integrate m services into n editors. Direct integration between editors and services (Figure 1a) would require $n \cdot m$ integrations to be implemented. To reduce this number, I propose to create an independent Real-Time Document Synchronization and Service Broker (ReDSyS) component, as shown in Figure 1, which complies with the practices **P1-P3**, and which integrates with each service and editor exactly once in a client-server manner (ReDSyS being the server). In this way we only need $n+m$ integrations. The ReDSyS server API is expected to be stable (requirement P2) hence any upgrades of editors/services may require adjustments in the integration only on the editor/service part.

Section 3.1 describes in more detail the ReDSyS component and how it integrates with editors and services. Section 3.2 presents the shared document model and shows how editing changes are represented. In section 3.3, I discuss management of change issues that can appear when integrating time consuming services as well as make explicit some requirements for reactive services. Section 3.4 presents the interaction model between users, editors and services.




Fig. 1. Direct integrations between editors and environments vs indirect integration through ReDSyS component.

3.1 The Real-Time Document Synchronization and Service Broker

The ReDSyS component has two main responsibilities:

1. provide a way for editors to trigger events to all or some subset of services,
2. allow editors and services to independently add/remove meta-data or text to the shared document in real-time.

The first responsibility is typical for broker components. In our case, it makes it possible for editors to request autocompletion suggestions (from all services) or request the type of a symbol (from e.g. Twelf services for LF documents [Pfe91]).

The second responsibility enables services to run independently, distributed on different systems and parallel to the editing process. So while the user is typing, a service might decide to start processing and integrate the results back into the document when finished. Whenever a user changes the document, services get notified by the ReDSyS component and each service has the freedom to decide whether to interrupt current processing (if available) or not. Thanks to real-time document editing solutions, it is quite often possible to automatically merge service results computed on older versions of the document into the current version.

To understand the interaction between editors, services and ReDSyS better, let me describe how $\text{\texttt{\textit{STeX}}}$ [Koh04] editing in Eclipse[Ecl] can be integrated with a term spotting and an autocompletion service. You can follow the communication between components in Figure 2.

The first step is to install the “ $\text{\texttt{\textit{STeX}}}$ -padconnector” plugin into Eclipse which, upon opening an $\text{\texttt{\textit{STeX}}}$ file, uploads it to the ReDSyS component and opens it in a typical Eclipse editing window. The opened document is, in fact, the shared document. The ReDSyS architecture takes care of starting the semantic services.

The “ $\text{\texttt{\textit{STeX}}}$ -padconnector” plugin is programmed so that, when the user presses `ctrl+space`, it synchronously notifies (i.e. waits for the result) the ReDSyS architecture that an event with a predefined URI “`autocomplete.stex`” took place and passes the current cursor coordinates as parameters. The autocomplete service catches the event and based on the parameters decides on autocompletion suggestions. The editor receives autocompletion suggestions from ReDSyS and displays them.

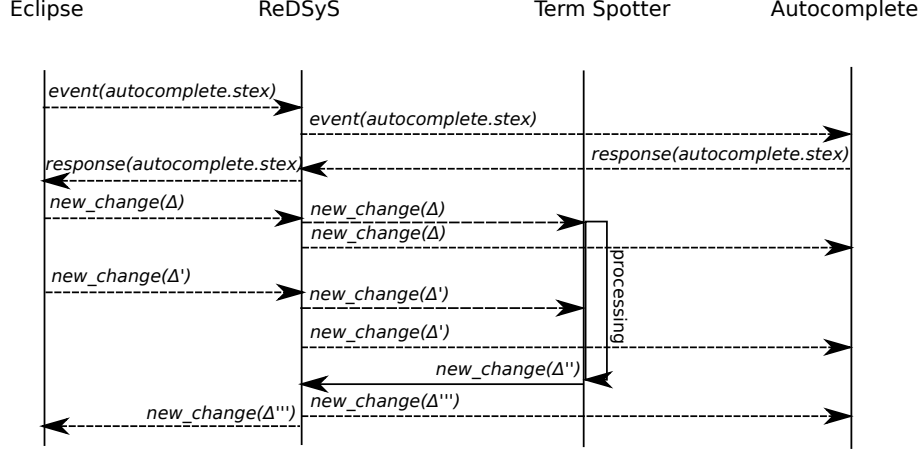


Fig. 2. Communication diagram among architecture components after an autocomplete requests followed by term spotting.

Whenever the user changes the document, a changeset (or diff) is computed and sent to ReDSyS. This passes on the changeset to all the other services so that they have an up-to-date version of the document. Let us suppose that the Term Spotter service decides to start a relatively complex NLP processing task for of the new version (e.g. version 40) of the document. While processing, some other changeset comes to the Term Spotter service but it decides not to cancel the NLP task. The shared document has now version 41 but the NLP task has computed a changeset (Δ'') which highlights new found terms based on document version 40. The Term Spotter service sends change Δ'' to ReDSyS also including information about the document version on which the changeset is based on (i.e. 40). The ReDSyS component tries to merge the changes and if it succeeds, sends a merged changeset to all the other components.

3.2 Document Model and Changesets

In this section I introduce the shared document model as it is important to understand what kind of information is shared among services and how. This document model is the same as that of the Etherpad-lite system [Eth].

We define a finite alphabet $A = \{\alpha_1, \alpha_2, \dots, \alpha_m\}$. A string $s_A = c_1c_2\dots c_n$ is a finite sequence of characters from alphabet A and length $len(s_A) = n$ is defined as the number of characters in that sequence. Let Str_A be the set of strings over alphabet A . For the sake of simplicity we will omit alphabet A in notations when it can be unambiguously inferred from the context.

An attribute pool $AP = \{(id, (key, val)) \in \mathbb{N} \times (Str \times Str) \mid id - \text{is unique}\}$ is a set of key-value pairs which also have a unique id assigned to each of them. This defines a function $attmap : \mathbb{N} \rightarrow (Str \times Str)$ returning the key-value pair associated to a certain id.

A document $d = (text, attmap, att)$ where $text \in Str$ represents the text in that document, $attmap \in \mathbb{N} \rightarrow (Str \times Str)$ is the function associated to an attribute pool and $att : \mathbb{N} \rightarrow Set(\mathbb{N})$ with $att(i) = S$, $\|S\| < \infty$ specifies which set of attributes are assigned to character i , $0 \leq i < len(t)$ in document's text.

A change operation $d = (op, S, len, t) \in \{+, -, =\} \times Set(\mathbb{N}) \times \mathbb{N} \times Str$ either

1. inserts (op="+") text t of length len and applies attributes in S to each of the inserted characters, or
2. deletes (op="-") len characters (and their attributes) from a text ($t = ""$ and $S = \emptyset$), or
3. leaves unchanged (op="=") len characters but applies attributes in S (if $S \neq \emptyset$) to each of them.

Editing changes inside a document are represented using lists of change operations $o = \{op_1 op_2 \dots op_k\}$ where op_j are change operations such that no two consecutive operations have the same type and attribute sets (otherwise we can join them together).

A changeset is defined as $c = (l, l', attmap, o)$, where l is the length of the text before the change, l' is the length of the text after the change, $attmap$ contains the updated attribute pool (only new elements allowed) and a sequence of change operations to be applied on the text.

To provide a better intuition for these notions, consider the document

$$d = \left(\text{"Math is great"}, \left\{ \begin{array}{l} 0 \rightarrow (\text{"bold"}, \text{"true"}) \\ 1 \rightarrow (\text{"author"}, \text{"p1"}) \\ 2 \rightarrow (\text{"author"}, \text{"p2"}) \end{array} \right\}, \left\{ \begin{array}{l} 0 \leq i \leq 4, \{1\} \\ 5 \leq i < 7, \{0, 2\} \\ 7 \leq i \leq 12, \{2\} \end{array} \right\} \right) \quad (1)$$

which says that word "Math" was authored by "p1", the rest of the text by "p2" and the word "is" is bold. Now consider a changeset

$$c = \left(13, 12, \left\{ \begin{array}{l} 0 \rightarrow (\text{"bold"}, \text{"true"}) \\ 1 \rightarrow (\text{"author"}, \text{"p1"}) \\ 2 \rightarrow (\text{"author"}, \text{"p2"}) \\ 3 \rightarrow (\text{"author"}, \text{"p3"}) \end{array} \right\}, \left\{ \begin{array}{l} (\text{" = "}, \{3\}, 1, \text{"p1"}), \\ (\text{" - "}, \emptyset, 3, \text{"p1"}), \\ (\text{" + "}, \emptyset, 2, \text{"MKM"}), \\ (\text{" = "}, \emptyset, 9, \text{"p1"}), \end{array} \right\} \right) \quad (2)$$

which when applied to d , would change the text to "MKM is great", the word "MKM" would have no author and the rest remains unchanged.

3.3 Time Consuming vs Reactive Services

The ReDSyS architecture can support both time consuming services (e.g. NLP tasks) that should not hinder the user from further editing of the document, as well as services that need to give user the impression that the service is running natively inside the editor (e.g. syntax highlighting).

As described in the communication workflow in the previous section, time consuming services can start processing at any given point in time (e.g. version 40) and integrate their results automatically (e.g. at version 50) by creating and sending a changeset (call it X) based on the version when processing started (i.e. 40). The ReDSyS component has some merging strategies to integrate such changes but they will fail if the changes done in between i.e. from version 41-50 overlap with areas in changeset X.

This default behavior can be improved in a number of ways. First, a time consuming service is still notified of changes done to the document even while it is processing. Hence, it can interrupt the processing or restart it if the document was changed in the area currently under processing. In this way, computing power to finish the processing (and then realize that it cannot be integrated) is not lost. The second solution is to try to incorporate incoming changes while or at the end of processing and ultimately create a changeset based on the latest version of the document. Hence a time consuming service is responsible for it's own management of change.

To give users the feeling that reactive services run natively in the editor, they need to be very optimized both in speed and in the size of the changesets they produce. These services might need to run at a rate of 20 times a second in order to accommodate several users editing in the same time. Hence it is very important that reactive services can cache results and start processing a document without reading all of it. Also the changesets that reactive services produce should be small and only change areas of the document that really need changing. For example, a bad syntax highlighting service that creates a changeset recoloring the whole document (and not only the parts the need recoloring) could invalidate the processing of all the time consuming services.

3.4 User Interaction Model

In the current framework, services and editors no longer integrate directly but they still need to interact, e.g. services might need to ask the user to disambiguate a mathematical term. Such interactions must be standardized so that all editors ask services to perform a certain action (e.g. autocomplete) in the same way.

Every type of interaction between users, editors and services has a predefined URI. An example of such URIs is "autocomplete.stex" which, when broadcasted by the ReDSyS component to the services, expects them to return \LaTeX based autocompletion suggestions and then displayed to the user. Another example is "contextmenu.spotter_plugin.10". This URI can be used inside an attribute ("ui","contextmenu.spotter_plugin.10") which, just like the ("bold","true") attribute, can be applied to some part of the text in the shared document model presented in section 3.2.

When the editor sees attributes having key "ui" and value prefixed with "contextmenu.", it knows to display a context menu when the text having this attribute is right-clicked. The menu items in the context menu are fetched from the "spotter_plugin" component (i.e. the Term Spotter) and "10" is passed as an argument to identify which context menu should be displayed.

It is the editor that is responsible to understand interaction URIs and act accordingly. That is why it is important to define interactions in a general and reusable manner so that many services can take advantage of it. Currently, the set of predefined interaction URIs is relatively small and fits the use cases presented in the implementation section. However, it certainly needs to be revised and extended to fit a more general range of interactions.

4 Architecture Implementation

To validate the presented architecture and to prove its applicability to a wide range of semantic services and editors, I chose to integrate four \LaTeX semantic services into two editors. The semantic services were picked to address different integration issues and are described in more detail in Section 4.1. Section 4.2 gives a glimpse into the APIs that need to be used in order to create a service. Finally, section 4.3 discusses extensibility and reuse issues of the architecture.

The editors I used to validate the architecture are: Eclipse (desktop based editor) and Etherpad’s Web Client (web-based editor). To support the Eclipse based editor I implemented my own synchronization library with the ReDSyS component called `jeasysync2`. The real-time document sharing platform (ReDSyS) used in my implementation is an extension of the Etherpad-lite system and can be found at <https://github.com/jucovschi/etherpad-lite/tree/mkm>.

4.1 Implemented Semantic Services

My architecture can be seen as an enabler for user-editor-service type of interactions and hence this is the part which needs most testing. Let us consider the service of semantic syntax highlighting. The user-editor-service interaction consists in the service being able to change the color of text. Testing a more complex service, requiring coloring parts of text does not make sense because the additional service complexity is independent of the presented architecture. Hence I chose services testing different aspects of user-editor-service interaction, namely:

\LaTeX semantic syntax highlighter colors \LaTeX code based on its semantic meaning. This is a service which cannot be implemented using regular expressions — the main tool for syntax highlighting in many editors. Hence I implemented it as a service and integrated it in editors via ReDSyS. Even though it only needs to highlight text, it has to do that more often than most other services hence it helps benchmarking the user-editor-service interaction speed.

Term Spotter is a NLP based service which tries to spot mathematical terms inside a document. The interaction with the user is very similar to that of spell checking, namely, spotted mathematical terms are underlined while the user is typing. The user can then choose to add semantic references to spotted terms. This is an example of a service with heavier server side part and helps us test how service results are automatically integrated (if possible) into newer document versions.

TermRef Hider and Transclusion services are examples of advanced editing features, one hiding parts of annotations from the user and other showing referenced text instead of references. As both services showcase very important results of using proposed architecture I address them in more detail in the next section.

Support for Advanced Editing Features

Inline annotated documents are very hard to author because they contain additional implicit knowledge that 1. is redundant to the author as she already knows it and 2. hinders a clear reading experience. Stand-off or parallel annotations solve these problems but they require the use of special editing environments every time a small update needs to be performed. Failing to do so may invalidate existing annotations. Generally MKM systems only support inline annotations on documents that are editable by users.

In Figures 3 and 4 compare a mathematical document to its semantically annotated version. The difference between the readability of these documents is quite obvious even though all we did was to annotate three terms (using `\termref` macros) and do four transclusions (using `\STRlabel` and `\STRcopy` macros).

```

1The gravitational potential energy of a system of masses $m_1$ and $M_2$
  at a distance $r$ using gravitational constant $G$ is
3\begin{equation}
  U = -G\frac{m_1M_2}{r}+K
5\end{equation}
  where $K$ is the constant of integration. Choosing the convention that $K=0$
7makes calculations simpler, albeit at the cost of making $U$ negative.

```

Fig. 3. Conventional mathematical document

```

1The \termref{cd=physics-energy, name=grav-potential}{gravitational potential energy}
  of a system of masses \STRlabel{m1}{m_1} \STRcopy{m1} and \STRlabel{m2}{M_2}
3\STRcopy{m2} at a distance \STRlabel{r}{r} \STRcopy{r} using
  \termref{cd=physics-constants, name=grav-constant}{gravitational constant}
5\STRlabel{G}{G} \STRcopy{G} is \STRlabel{U}{U} \STRcopy{U}
  \begin{equation}
7  \STRcopy{U} = -\STRcopy{G}\frac{\STRcopy{m1}\STRcopy{m2}}{\STRcopy{r}}+
    \STRcopy{K}
9\end{equation}
  where \STRcopy{K} is the \termref{cd=physics-constants, name=integration}{constant
11of integration}. Choosing the convention that \STRcopy{K}=0 makes calculations
  simpler, albeit at the cost of making \STRcopy{U} negative.

```

Fig. 4. Semantically annotated mathematical document

To make the text in figure 4 look as readable as the one in 3, we need to support 2 features, namely: inline folding and transclusion. Using inline folding, one could collapse a whole `\termref` to show only the text in its second argument. The transclusion feature would then replace `\STRcopy` references with the text in the `\STRlabel`.

The inline folding or transclusion features are not supported by most editing environments used to author MKM formats like \LaTeX , Mizar[UB06] or LF[Pfe91]. Adding these features directly in each of the authoring environments requires a lot of initial development effort and incurs high maintenance costs when the editor evolves.

4.2 Libraries and APIs

Currently one can create new services for my architecture using either JavaScript or Java programming languages. JavaScript services are implemented in Etherpad-lite’s native plugin system. Java services should use the jeasysync2 library. In both cases services must implement the following interface:

```

void init(Changeset initialText, AttributePool pool);
2 void update(Changeset lastChangeset, AttributePool newPool, ChangesetAcceptor csAcceptor);

```

The init method is called when initializing the service. The first parameter is the changeset which, if applied to an empty text, generates the current document (note that attributes are included as well).

The update method notifies the service of new updates. This is where the service should decide whether to start/restart processing. The update function is called asynchronously in separate threads so special care should be taken not to run into race conditions. The ChangesetAcceptor callback allows the service to send changesets back to the ReDSyS component when processing is finished.

Creating changesets is easily done through a utility class called ChangesetBuilder with the methods:

```

void keep(int noChars, AttributeList attribs);
2 void insert(String text, AttributeList attribs);
void remove(int noChars);

```

This class allows services to specify changes they want to perform in a sequential way e.g. keep the first 10 characters untouched, remove the next 5, insert text “Hello World” and apply attribute [“bold”, “true”] to it, keep the next 2 characters unchanged but apply attribute [“bold”, “”] to them etc. The ChangesetBuilder class will produce a correctly encoded changeset which can be then transmitted to ReDSyS.

The last best practice I would like to share is a simple and efficient algorithm of converting a list of changes of type “*apply attribute [key_i, value_i] from character begin_i to character end_i*” to a sequential list of changes suitable for the ChangesetBuilder. I used this algorithm (with minor changes) for all 4 services, hence might be of interest for future service developers.

- we create a list of “event” triples having signature (*type, i, attr*) where *type* is either “add” - to add attribute “attr” to the list of attributes applied to all following characters, or “remove” to remove attr from the list of attributes. Index *i* specifies at which position in the sequence a certain event should take place.

- for each rule of type “*apply attribute [key_i, value_i] from character begin_i to character end_i*” add event triples (“add”, *begin_i*, [*key_i*, *value_i*]) and (“remove”, *end_i* + 1, [*key_i*, *value_i*])
- sort the event list by the *i* values
- initialize an empty list of attributes called *currentAttrs* and set *lastPos* = 0
- iterate through the sorted event list and let (*type*, *i*, *attr*) be the current event
 - if *i* > *lastPos*, add sequential operation *keep(i – lastPos, currentAttrs)* and set *lastPos* = *i*.
 - if *type* = “add”, add *attr* to *currentAttrs*
 - if *type* = “remove”, remove *attr* from *currentAttrs*

This algorithm can be generalized to handle events which delete or insert text as well.

4.3 Evaluation of Integration Costs

The aim of the presented architecture is to minimize integration costs and hence in this section I want to evaluate what we gain by using it.

Costs for Integrating Custom Editors / Services

Both editors and services need to implement the following functionality

RE1. Connect to the ReDSyS component,

RE2. Implement document model and changeset synchronization mechanisms. Both RE1 and RE2 can be reused from other already integrated editors/services. If no such implementations exist for a certain programming language, my own experience shows that one needs to invest about one day for implementing RE1 and about two weeks for RE2. Implementing RE2 requires mostly code porting skills (for about 2k lines of code) i.e. good understanding of particularities of programming languages but does not require deep understanding of the algorithms themselves. Additionally, unit tests help a lot finding and fixing bugs. Currently, RE1 and RE2 are available for Java and JavaScript languages using *jeasysync2*¹ and *easysync2*² libraries.

Editors need to additionally implement styling and user interaction mechanisms based on attributes in the shared document model. Depending on the editing environment, this might take several more days.

Integration of services requires implementation of the interaction of the service with the shared document and with the user. To integrate each service described in section 4, I needed 50-100 lines of code.

Requirements for Integrating a Custom Interaction

To add new user interactions, one has to choose a unique URI to identify the interaction and then use this URI from an editor (to trigger events) or

¹ <https://github.com/jucovschi/jeasysync2>

² <https://github.com/Pita/etherpad-lite/static/js/Changeset.js>

from services by using it inside an attribute. The biggest cost associated with adding/extending the set of interactions is that of propagating it to all already integrated editors. Versioning and change management of interaction URIs is an open issue.

5 Conclusion & Future Work

The extent to which editing environments could support the authoring process of MKM documents is far from being reached. There are lots useful authoring services which, if integrated in editing environments, would make authoring of MKM documents easier to learn, more efficient and less error-prone.

Many MKM editing services are available only in certain editing environments but even more services live solely in the wish-lists of MKM authors. An important reason for it is that creating and maintaining such integrations is very expensive. This paper suggests that integration of editors with MKM authoring services can be done in an efficient way. Namely, a service showing the type of LF symbols needs to be integrated with the ReDSyS component once and then be used in a ever-growing list of editors like Eclipse, jEdit or even web-editors. Conversely, once an editor (e.g. \TeX macs) integrates with the ReDSyS component, it would be able to provide the user with all the services already integrated with the ReDSyS component.

The implementation part of this paper allowed me to test my ideas and I found out that integrating a service into all already integrated editors can take as little as 3-4 hours of work. In case that no communication and synchronization libraries are available for a certain language, one can implement them in about 2-3 weeks,

While the focus of the current paper is to reduce the costs for integrating m services into n editors, the suggested solution also:

1. makes is possible for services which need longer processing times to run in the background without interrupting the user's authoring experience,
2. allows services to be implemented in any convenient programming language or framework and even be distributed on different hardware, and
3. extends the typical plain-text document model with attributes which provide a very convenient storage for layers of semantic information inferred by services.

Future research directions include development and integration of new MKM services into editors, extending the list of programming languages which can connect to the ReDSyS component, and integrating additional editors into the proposed architecture.

References

- [ALW06] D Aspinall, C Luth, and B Wolff. "Assisted Proof Document Authoring". In: *Lecture Notes in Computer Science* 3863 (2006), pp. 65–80.

- DOI: 10.1007/11618027_5. URL: <http://www.springerlink.com/index/fq4068582k604115.pdf>.
- [ALW07] David Aspinall, Christoph Lüth, and Daniel Winterstein. “A Framework for Interactive Proof”. In: *Calculemus/MKM*. Ed. by Manuel Kauers et al. LNAI 4573. Springer Verlag, 2007, pp. 161–175. ISBN: 978-3-540-73083-5.
- [Ecl] *Eclipse: An open development platform*. seen December 2011. URL: <http://www.eclipse.org/>.
- [Eth] *Etherpad lite: real time collaborative editor*. web page at <https://github.com/Pita/etherpad-lite/raw/master/doc/easysync/easysync-notes.pdf>. seen Feb 2012. URL: <https://github.com/Pita/etherpad-lite/raw/master/doc/easysync/easysync-notes.pdf>.
- [Koh04] Michael Kohlhase. “Semantic Markup for \TeX / \LaTeX ”. In: *Mathematical User Interfaces*. Ed. by Paul Libbrecht. 2004. URL: <http://www.activemath.org/~paul/MathUI04>.
- [Pfe91] Frank Pfenning. “Logic Programming in the LF Logical Framework”. In: *Logical Frameworks*. Ed. by Gérard P. Huet and Gordon D. Plotkin. Cambridge University Press, 1991.
- [TN92] I Thomas and B A Nejme. “Definitions of tool integration for environments”. In: *IEEE Software* 9.2 (1992), pp. 29–35. ISSN: 07407459. DOI: 10.1109/52.120599. URL: <http://doi.ieeecomputersociety.org/10.1109/52.120599>.
- [UB06] Josef Urban and Grzegorz Bancerek. “Presenting and Explaining Mizar”. In: *Proceedings of the International Workshop “User Interfaces for Theorem Provers” 2006 (UITP’06)*. Ed. by Serge Autexier and Christoph Benzmüller. Seattle, USA, 2006, pp. 97–108.
- [Was90] Anthony L Wasserman. “Tool integration in software engineering environments”. In: *Development* 1.6 (1990), pp. 137–149. ISSN: 02686961. DOI: 10.1007/3-540-53452-0_38. URL: <http://www.springerlink.com/content/p582q2n825k87n15>.
- [Wen10] Makarius Wenzel. “Asynchronous Proof Processing with Isabelle/Scala and Isabelle/jEdit”. In: *User Interfaces for Theorem Provers (UITP 2010), FLOC 2010 Satellite Workshop*. Ed. by C. Sacerdoti Coen and D. Aspinall. ENTCS. Elsevier, 2010. URL: <http://www.lri.fr/~wenzel/papers/async-isabelle-scala.pdf>.
- [Wic04] Mike Wicks. “Tool Integration in Software Engineering : The State of the Art in 2004”. In: *Integration The Vlsi Journal* August (2004), pp. 1–26. URL: <http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.59.1969>.